arXiv:cs/0106008v1 [cs.PL] 7 Jun 2001

Computing Functional and Relational Box Consistency by Structured Propagation in Atomic Constraint Systems *

M.H. van Emden
Constraints Group, CWI and Aerospace Faculty, TU Delft
On leave from Computer Science Dept, University of Victoria

Abstract

Box consistency has been observed to yield exponentially better performance than chaotic constraint propagation in the interval constraint system obtained by decomposing the original expression into primitive constraints. The claim was made that the improvement is due to avoiding decomposition. In this paper we argue that the improvement is due to replacing chaotic iteration by a more structured alternative.

To this end we distinguish the existing notion of box consistency from *relational* box consistency. We show that from a computational point of view it is important to maintain the functional structure in constraint systems that are associated with a system of equations. So far, it has only been considered computationally important that constraint propagation be fair. With the additional structure of functional constraint systems, one can define and implement computationally effective, structured, truncated constraint propagations. The existing algorithm for box consistency is one such. Our results suggest that there are others worth investigating.

1 Introduction

Systems of nonlinear equations where the unknowns are reals arise in specialized applications such as the study of chemical equilibrium and in robot kinematics. A general class of applications of systems of nonlinear equations arises when optimizing a function of n variables with multiple local minima. A common optimization method is to set the n partial derivatives to zero, and solve the resulting set of n equations, which are in general nonlinear. Thus, nonlinear equations occur widely in mathematical modeling.

Until recently, only Newton's method was available for solving such a system. This method is good at refining sufficiently good estimates of a solution. Trying to use it otherwise is a hit-and-miss affair. The situation was greatly improved with the advent of interval arithmetic [8, 9].

*Presented at the Sixth Annual Workshop of the ERCIM Working Group on Constraints, June 18-20, 2001, Charles University, Prague. This document is report DCS-266-IR, Department of Computer Science, University of Victoria, Victoria, BC, Canada.

A remarkable subsequent development was BNR Prolog [4] (later referred to as "CLP(Intervals)"), which introduced what came to be called *interval constraints*. This method can be regarded as an adaptation of the CHIP system [6]. This constraint processing system associates each variable with a finite set of possible values. Instead, in BNR Prolog, the set of possible values is an interval of reals. BNR Prolog adopted from CHIP an instance of a constraint propagation algorithm that turned out to be an instance of Apt's Generic Chaotic Iteration algorithm ("GCI" in the sequel) [1].

To solve nonlinear systems with CLP(Intervals), one has the advantage of not needing derivatives. The equations are decomposed into primitive constraints. The resulting constraint system is then subjected to a constraint propagation algorithm (the *pruning step*). The resulting box is split (the *branching step*), whereupon the same is done recursively to the results of the split. By ensuring that pruning preserves completeness (that is, does not remove any solutions), one ensures that the solving algorithm generates a sequence of boxes that contain all solutions.

Benhamou, McAllester, and Van Hentenryck [2] showed that this simplicity comes at a cost: on the Broyden Banded Function, a widely used benchmark, exponentially increasing computation time is needed when one only uses constraint propagation for pruning. This is not surprising because of the presence of branching.

By contrast [2] describes `Newton`, an algorithm that does not require branching in this benchmark and only exhibits linearly increasing computation time. This remarkable improvement was based on the novel notions of *box consistency* and of the *pseudo-zeros* used to characterize maximally box-consistent sets.

`Newton` avoided branching on this particular benchmark because box-consistency achieved stronger pruning than the constraint propagation used in CLP(Intervals). However, box consistency was a step back in the sense of only using functional interval arithmetic, which contracts only the interval for the function value, rather than using the propagation of CLP(Intervals), which has the potential of contracting all intervals involved in a relation. Logically, the next improvement was to be one that combined the advantages of box consistency with those of the relational interval arithmetic used in CLP(Intervals). This step was taken by Benhamou, Goualard, Granville, and Puget with their HC4 algorithm [3].

In this paper, we explore other ideas for improving the use of relational interval arithmetic. We describe the use of *probing* an interval constraint system to improve the bounds obtained by a single constraint propagation. We show that this leads to a relational form of box consistency, which is stronger than the original notion of box consistency, which we call *functional* box consistency. We show that functional and relational box consistency are but two extremes of a spectrum defined by ways of structuring and truncating the iteration in constraint propagation. One of these ways can be regarded as the simulation of the evaluation of an expression in interval arithmetic.

Disclaimer Many basic definitions and results need to be covered here. In most cases, no attempt at attribution will be made. In the interest of mutual compatibility, some definitions are modified. As a result attribution might not be welcomed, yet no novelty is involved. Possible novelties are simulation of interval arithmetic by constraint propagation, relational box consistency and its

computation by probing, and the identification of alternatives to full constraint propagation based on structured rather than chaotic iteration.

2 Constraints and equations

2.1 Constraint systems

Definition 1 *A constraint system has the following attributes.*
(1) A set $\{T_1, \ldots, T_n\}$ *of sets called* types.
(2) A set $\{x_1, \ldots, x_n\}$ *of variables where* x_i *is of type* T_i *for* $i \in \{1, \ldots, n\}$.
(3) A set $\{A_1, \ldots, A_m\}$ *of constraints where* A_i *is an atomic formula of first-order predicate logic.* $\{x_1, \ldots, x_n\}$ *is the set of all variables occurring in* $\{A_1, \ldots, A_m\}$. *For* $i = 1, \ldots, m$, $d_i \subseteq \{1, \ldots, n\}$ *is such that* $\{x_j \mid j \in d_i\}$ *is the* set of variables occurring in A_i.
(4) A state, *which is* $D_1 \times \cdots \times D_n$, *where, for* $i \in \{1, \ldots, n\}$, $D_i \subseteq T_i$. *We say that, in this state of the constraint system,* D_i *is the* domain *of* x_i.
(5) An initial state, *which is a state.*

Definition 2 *The* relation associated with *a function or operation* $f : \mathcal{R}^k \to \mathcal{R}$ *is* $\{\langle x_0, \ldots, x_k\rangle \mid x_0 = f(x_1, \ldots, x_k)\}$, *for* $k = 0, 1, \ldots$

We will say that a function or operation is *admissible* if the contraction operator of the associated relation is efficiently computable. This means roughly that it can be computed without iteration. Admissible operations include addition, subtraction, multiplication, division, maximum, absolute value, power for all integer exponents, exp, log, and the trigonometric functions.

Proposition 1 *Let* r *be the relation associated with an admissible* $f : \mathcal{R}^k \to \mathcal{R}$. *Let the intervals associated with* $x_0, x_1, \ldots, x_k$ *be* $[-\infty, +\infty], I_1, \ldots, I_k$, *respectively. Then the intervals resulting from applying the contraction operator of* r *are, respectively, the intervals* $f'(I_1, \ldots, I_k), I_1, \ldots, I_k$, *where* f' *is the canonical set extension of* f.

Definition 3 *An* equation *is* $E = 0$, *where* E *is an expression of type real containing only real variables. An* equation system *consists of a set* X *of real variables and a set of equations containing no variables other than those in* X.

In the conventional way, we consider expressions as trees. The leaf nodes are constants or variables; the nonleaf nodes are the operation symbols.

A distinguishing feature of CLP(Intervals) is that it decomposes equations, or other composite expressions, into primitive constraints. These primitive constraints are the relational versions of the building blocks of expressions, which are admissible functions.

It is this decomposition, described in the definition below, that has been identified in [2] as the cause for the observed slowness of CLP(Intervals).

Definition 4 *The* constraint system $\mathcal{C}$ associated *with an equation system* $\mathcal{E}$ *depends on a one-one correspondence between the non-leaf nodes of the trees in* $\mathcal{E}$ *and a set of variables that is disjoint from the variables in* $\mathcal{E}$, *which is defined as follows.*

Each of the variables in $\mathcal{E}$ also occurs in $\mathcal{C}$, where it is called a "primary variable". To each non-leaf node of an expression in $\mathcal{E}$ there corresponds a variable in $\mathcal{C}$ that does not occur in $\mathcal{E}$ and is called "auxiliary variable".

The constraints in $\mathcal{C}$ are determined as follows. For every non-leaf node n (which is an operation f), with children $n_1, \ldots, n_k$, of a tree in $\mathcal{E}$, there is an atomic formula in $\mathcal{C}$ with variables $x_0, x_1, \ldots, x_k$. The predicate in the atomic formula denotes the relation associated with f. This formula is a functional atom *in $\mathcal{C}$. The variables $x_1, \ldots, x_k$ are the* input variables *of the atom; x_0 is its* output variable.

In addition there is in $\mathcal{C}$, for every root of a tree in $\mathcal{E}$ with corresponding variable v, an atomic formula $v = 0$. This is a relational atom *in $\mathcal{C}$.*

Definition 5 (Floating-point numbers, intervals) *A floating-point number is any element of $F \cup \{-\infty, +\infty\}$, where F is a finite set of reals that includes 0. If x is a finite floating-point number, then x^- (x^+) is the greatest (smallest) floating-point number smaller (greater) than x. In addition, $-\infty^- = -\infty$, $-\infty^+ = -M$, $+\infty^- = M$, and $+\infty^+ = +\infty$, where M is the greatest finite floating-point number.*

A floating-point interval *is a closed connected set of reals, where the bounds, in so far as they exist, are floating-point numbers. When we write "interval" without qualification, we mean floating-point interval.*

An interval that does not properly contain an interval is called canonical.

A box *is a cartesian product of floating-point intervals.*

Thus canonical intervals are non-empty sets of reals. They may have positive width and they may have zero width. Examples are $[a^-, a]$, $[a, a^+]$, and $[a, a]$, where a is a finite float-point number. For any real, there is a unique smallest canonical floating-point interval containing it.

In this paper we consider *interval constraint systems*, which are constraint systems where the types are all equal to $\mathcal{R}$ and where the domains are intervals.

3 Propagation

Apt's Generic Chaotic Iteration algorithm (GCI) [1] is of an astonishing simplicity and wide applicability. The elegance of CLP(Intervals), noted in [2], is due in part to the fact that its constraint propagation is an instance of GCI.

GCI maintains a pool of operators that still need to be applied. The attraction of GCI is that it does not specify any order among these applications. In this section we consider orders of application that are computationally effective for constraint systems that are associated with equation systems. We formalize application order by means of "traces", as defined below.

Definition 6 *A* trace *for a given constraint system $\mathcal{C}$ with attributes as in Definition 1 has the following components:*

(1) An index sequence *t, which is an infinite sequence with elements in $\{1, \ldots, m\}$.*

(2) A sequence-of-atoms *of which the i-th element is the atom A_{t_i} in $\mathcal{C}$, for $i = 0, 1, \ldots$*

(3) A sequence-of-constraints *of which the i-th element is the relation defined by the atom A_{t_i} in $\mathcal{C}$, for $i = 0, 1, \ldots$*

(4) A sequence-of-contraction operators *of which the i-th element is the contraction operator τ_{t_i} defined by the atom A_{t_i} in $\mathcal{C}$, for $i = 0, 1, \ldots$*

(5) A sequence-of-boxes *U of which the i-th element is the initial box of $\mathcal{C}$ if $i = 0$ and is $\tau_{t_{i-1}}(U_{i-1})$ if $i > 0$.*

As we are primarily interested in the sequence of boxes, we think of the sequence of contraction operators as "activations" of the corresponding constraints. We think of the elements of t as "selecting" a constraint to be activated.

The following proposition is based on the fact that the contraction operators are monotone nonincreasing and idempotent and that there are finitely many domains.

Proposition 2 *See [10, 1].*
For any interval constraint system with box B as initial state we have:
(1) The sequence of boxes has a limit for every trace.
(2) The greatest lower bound of these limits is also a limit of a trace.
(3) All traces in which the index sequence is fair have the same limit. This limit equals the greatest common fixpoint of $\tau_1, \ldots, \tau_m$ that is less than B.
(4) This fixpoint is uniquely determined by the constraint system. It is computed by a suitable instance of GCI.

Definition 7 *A constraint system is* failed *(*non-failed*) if its fixpoint is empty (non-empty).*

A failed constraint system has no solutions. A non-failed constraint system may, but need not, have solutions.

Proposition 3 *The fixpoint of a constraint system contains all its solutions.*

Definition 8 *A segment of a trace is* functional *if*

(1) The sequence-of-atoms only contains functional atoms.

(2) For every atom, any input variable that is an auxiliary variable has occurred as output variable earlier in the segment.

(3) No atom occurs more than once.

A segment of a trace is inverse functional *if every occurrence of a variable as an output variable has been preceded by an occurrence of that variable as input variable or as variable in an activation of an equality constraint.*

The following proposition shows that a trace of a constraint system can simulate the evaluation of an expression.

Proposition 4 *Let $\mathcal{C}$ be an interval constraint system with attributes as in Definition 1 that is associated with an equation system $\mathcal{E}$ containing an expression E. Let box B be the initial state of $\mathcal{C}$ such that all its projections corresponding to auxiliary variables are $[-\infty, +\infty]$. Let $d \subseteq \{1, \ldots, n\}$ be such that $\{x_j \mid j \in d\}$ is the set of variables in E. Let x_i be the variable in $\mathcal{C}$ that corresponds to the root of E. Let $\mathcal{T}$ be a functional initial segment of a trace.*

The i-th projection of the last of the sequence-of-boxes of $\mathcal{T}$ is the same interval as the one obtained when E is evaluated in interval arithmetic with $\pi_j(B)$ as the interval substituted for x_j in E, for all $j \in d$.

Consequences of Proposition 4 GCI leaves open the possibility of activating an operator that has no effect; that is, when this activation does not result in any domain reduction. In any trace containing an operator activation without effect, the number of steps to convergence can be decreased by removing it.

One possible heuristic to avoid vacuous activation is to require a trace to be a *two-phase iteration*: to consist of a repetition of cycles consisting of the following two items: (1) a functional segment and (2) an inverse-functional segment.

It may be conjectured that such a trace is optimal in the sense of having a shortest pre-convergence initial segment. It is to be expected that such a trace is subject to severely diminishing returns in the sense of activations resulting in domain reduction.

The first functional segment likely has the greatest effect. It corresponds to the interval-arithmetical special case of constraint propagation. While one can probably construct examples to the contrary, most of the remaining reduction is typically effected in the first inverse-functional segment. For some purposes, to be discussed below, it is advisable to truncate constraint propagation before convergence has occurred. Promising truncations are: after the first functional segment or after the first cycle.

4 Box consistency

We first describe, adapted to the current setting, the box consistency notion of [2]. To distinguish it from the version to be described next, we call it "functional box consistency".

4.1 Functional box consistency

Definition 9 (Coordinate-wise functional box consistency operator) *Given an equation $E = 0$ with variables $x_1, \ldots, x_n$, and $i \in \{1, \ldots, n\}$. The i-th* coordinate-wise functional box consistency operator *replaces in $I_1 \times \cdots \times I_n$ the interval I_i by its intersection with the smallest interval containing $\{y \in \mathcal{R} \mid 0 \in E'\}$, where E' is the result of evaluating E in interval arithmetic with $I_1, \ldots, I_n$ substituted for the variables $x_1, \ldots, x_n$, except that the smallest floating-point interval containing y is substituted for x_i.*

The i-th coordinate-wise functional box consistency operator is completeness-preserving in the sense that it removes no solution to $E = 0$.

Definition 10 (Functional box consistency) *A box is* functionally box consistent *with respect to an equation if it is a common fixpoint of the n coordinate-wise functional box consistency operators associated with that equation.*

This is equivalent to the notion of box consistency introduced in [2]. By itself, box consistency is not interesting: for example, the empty box is box-consistent with respect to $\{x = 0\}$, which has plenty of zeros. What is lacking so far is a suitable notion of maximality. This was done in [2] by relating box consistency with leftmost and rightmost pseudo-zeros. It can also be done by means of fixpoint theory. Because each coordinate-wise functional box consistency operator is a monotonic one defined on the partially ordered set of boxes with set containment as partial order, it has a greatest fixpoint. By applying GCI one

obtains the greatest fixpoint common to all n coordinate-wise functional box consistency operators associated with a set of equations.

Definition 11 (Functional box consistency operator) *Given an equation $E = 0$ with variables $x_1, \ldots, x_n$. The* functional box consistency operator *is the mapping from a box B to the greatest fixpoint contained in B that is common to all n coordinate-wise functional box consistency operators associated with the equation.*

Definition 12 (Functional pseudo-zero) *Given an expression E with variables $x_1, \ldots, x_n$. An* i-th functional pseudo-zero *of E with respect to the intervals $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$ is a canonical interval I such that 0 is contained in the interval resulting from evaluating E in interval arithmetic with I substituted for x_i and I_j substituted for x_j, for all $j \in \{1, \ldots, n\} \setminus \{i\}$.*
An i-th functional pseudo-zero of an equation system *with respect to the intervals $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$ is a canonical interval that is a functional pseudo-zero of every one of its equations.*

Proposition 5 *Given an expression E with variables $x_1, \ldots, x_n$. If, for any $i \in \{1, \ldots, n\}$, no i-th functional pseudo-zero exists with respect to $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$, then $E = 0$ has no solution with $x_j \in I_j$, for all $j \in \{1, \ldots, n\} \setminus \{i\}$.*

Proposition 6 *The result of applying the functional box consistency operator of an equation system with at least one solution to a box B results in a box that has as i-th projection the least interval containing the i-th leftmost and rightmost functional pseudo-zeros with respect to $\pi_1(B), \ldots, \pi_n(B)$, for all $i \in \{1, \ldots, n\}$.*

4.2 Probing

Definition 13 *Let $\mathcal{C}$ be a constraint system.* Probing $\mathcal{C}$ with a constraint A *means determining whether the constraint system $\mathcal{C}'$, which is the result of adding A to $\mathcal{C}$, is failed or non-failed.*

Note that $\mathcal{C}$ does not change as the result of probing.

Proposition 7 (Monotonicity of probing) *If a non-failed constraint system $\mathcal{C}$ yields failure as a result of probing with $x \leq u_1$, then it also yields failure as a result of probing with $x \leq u_2$, where u_1 and u_2 are floating-point numbers such that $u_2 < u_1$.*

A similar fact holds for probing with $x \geq l_1$ or $x \geq l_2$, with $l_2 > l_1$. Also non-failure can be inferred from non-failure by means of suitably selected probes.

When probing a constraint system $\mathcal{C}$ containing a variable x with $x \leq a$ yields failure, then it has been proved that no solution has an x-component that is less than or equal to a. It does not follow that probing $\mathcal{C}$ with $x \geq a$ yields non-failure. It is quite common for $\mathcal{C}$ to be non-failed, yet not to have any solutions. This is sometimes discovered by probing $\mathcal{C}$ with both $x \leq a$ and $x \geq a$ and finding failure in both cases. While this property of probing is often effective, such an x and such an a cannot always be found.

As seen above, probing has a logic of its own. Here is an other example. It may be that a non-failed constraint system $\mathcal{C}$ containing variable x yields non-failure on probing with $x \geq a$ and that it yields non-failure on probing with $x \leq b$, where $a < b$, and yet yields failure on probing with $x \geq a \wedge x \leq b$. Conversely, non-failure on probing with $x \geq a \wedge x \leq b$ implies non-failure on probing with $x \geq a$; it also implies non-failure on probing with $x \leq b$.

Probing suggests a relational version of the functional pseudo-zero. As probing applies to any constraint system, not just to those derived from equations, we call them "pseudo-solutions".

4.3 Relational box consistency

Probing can be used to compute approximations to relational box consistency, a criterion similar to functional box consistency, but one that produces closer approximations to the set of solutions. Another way in which relational box consistency is interesting is that it applies to all interval constraint systems, not only to those that are derived from equations.

We proceed in analogy with Section 4.1. In analogy with the pseudo-zeros of [7] we have the following definition.

Definition 14 ((Canonical) pseudo-solution) *Let $\mathcal{C}$ be a non-failed interval constraint system with initial box $I_1 \times \cdots \times I_n$ and let $i \in \{1, \ldots, n\}$. Let $\mathcal{C}'$ be an interval constraint system with the same attributes as $\mathcal{C}$, except that interval I_i is changed to an interval y. If $\mathcal{C}'$ is non-failed, then y is an* i-th pseudo-solution of $\mathcal{C}$ with respect to $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$. *If y is a canonical interval, then it is an* i-th canonical pseudo-solution of $\mathcal{C}$ with respect to $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$.

Functional pseudo-zeros have been defined, following [2, 7] as being canonical intervals. Here it is useful that pseudo-solutions are not necessarily canonical.

Definition 15 (Coordinate-wise relational box consistency operator) *The i-th* coordinate-wise relational box consistency operator *of an interval constraint system $\mathcal{C}$ maps the initial box $I_1 \times \cdots \times I_n$ to one where the i-th projection ($i \in \{1, \ldots, n\}$) has been replaced by the least floating-point interval containing the union of the i-th canonical pseudo-solutions of $\mathcal{C}$ with respect to $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$.*

We introduced Definition 10 because it figures so prominently in the literature, following [2]. One could define a relational analog, but it is more useful to skip to the analogs of Definition 11 and Definition 12.

Definition 16 (Relational box consistency operator) *Given an interval constraint system $\mathcal{C}$. The* relational box consistency operator *of $\mathcal{C}$ maps its initial box $I_1 \times \cdots \times I_n$ to the greatest fixpoint contained in it that is common to all n coordinate-wise relational box consistency operators associated with $\mathcal{C}$.*

Definition 17 (Leftmost (Rightmost) pseudo-solution) *Let $\mathcal{C}$ be an interval constraint system with variables $x_1, \ldots, x_n$ and let $i \in \{1, \ldots, n\}$. An i-th*

pseudo-solution $[a, b]$ *of* $\mathcal{C}$ *is an* i-th leftmost pseudo-solution of $\mathcal{C}$ *if* $\mathcal{C}$ *has no solution* $\langle \xi_1, \ldots, \xi_n \rangle$ *with* ξ_i *in* $\{x \in \mathcal{R} \mid x \leq y \text{ for all } y \in [a, b]\}$[1].

Pseudo-solutions become less interesting the wider they are. For example, $[-\infty, +\infty]$ is a pseudo-solution, a leftmost, and a rightmost pseudo-solution for any interval constraint system. On the other hand, consider an i-th canonical pseudo-solution $[a, b]$ such that no i-th canonical pseudo-solution is to the left of it. As every real ξ_i that is the i-th component of a solution is contained in an i-th canonical pseudo-solution, $[a, b]$ is also an i-th leftmost pseudo-solution. Hence the following proposition.

Proposition 8 *Suppose the relational box consistency operator maps box* B *to* B'. *Then* B' *is equal to the smallest box containing the* i*-th leftmost and rightmost canonical pseudo-solutions, for* $i \in \{1, \ldots, n\}$.

The reason for having these different characterizations for the same object is that the one in terms of pseudo-zeros is convenient for a relational box-consistency algorithm; see section 5. The one in terms of fixpoints is convenient for comparing relational box consistency with functional box consistency.

Proposition 9 *Let* E *be an expression with* n *variables and let* $\mathcal{C}$ *be the interval constraint system associated with* $\{E = 0\}$. *The result of mapping a box* B *with the relational box consistency operator of* $\mathcal{C}$ *is contained in the result of mapping* B *with the functional box consistency operator of* $\{E = 0\}$.

Proof. Both mapping results are characterized as common greatest fixpoints. Hence it is sufficient to show that, with respect to $\{I_j \mid j \in \{1, \ldots, n\} \setminus \{i\}\}$, the result of the i-th relational box consistency operator is contained in that of the i-th functional box consistency operator.

This can be ascertained by considering a canonical interval $[a, b]$ for which $0 \notin f_i([a, b])$. Here f_i is a function from intervals to intervals such that $f_i(I)$ is the result of evaluating E in interval arithmetic with I_j substituted for x_j for $j \in \{1, \ldots, n\}$ except that I is substituted for x_i.

Such an interval $[a, b]$ lies outside the result of applying the i-th functional box consistency operator. One fair trace of $\mathcal{C}$ begins with a functional segment followed by activating the equality constraint. It yields the empty interval at that point, which ensures that the unique limit of any fair trace is empty. Hence $[a, b]$ also lies outside the result of the i-th relational box consistency operator.

5 Computing Box Consistency

As only maximal box-consistent boxes are of interest, we compute leftmost and rightmost pseudo-zeros.

5.1 Computing the leftmost functional pseudo-zero

Given an equation E with variables $x_1, \ldots, x_n$. The purpose of function ZERO1 described in this section is to compute the i-th functional leftmost pseudo-zero of E. It is an adaptation of **function** `LeftNarrow` in [7].

[1] It is simpler to say: "with $\xi_i \leq a$", but this does not yield a useful definition because ξ_i is a real and a is a floating-point number. As a result, a can be $-\infty$.

```
interval function ZERO1(a, b) {
    if (0 ∉ f([a, b])) return ∅;
    if ([a, b] is canonical) return [a, b];
    //[a, b] not canonical, so has midpoint
    m := midpoint of [a, b];
    I' := ZERO1(a, m);
    if (I' is empty) return ZERO1(m, b);
    return I';
}
```

Figure 1: A definition of a function to compute a functional leftmost pseudo-zero.

```
interval function ZERO2(a, b) {
    if ([a, b] is canonical) return [a, b];
    m := midpoint of [a, b];
    probe C with x_i ≤ m;
    if (result of probing is failure) {
        probe C with m ≤ x_i ∧ x_i ≤ b;
        if (result of probing is failure) return ∅;
        return ZERO2(m, b);
    }
    I' := ZERO2(a, m);
    if (I' is empty) return ZERO2(m, b);
    return I';
}
```

Figure 2: A definition of a function to compute the i-th canonical leftmost pseudo-solution.

Suppose function ZERO1 is first called with arguments a_0 and b_0 such that the i-th leftmost functional pseudo-zero, if it exists, is contained in $[a_0, b_0]$. It returns an interval $[a', b']$. If this interval is empty, the algorithm has detected that no solution occurs in $[a_0, b_0]$. Otherwise, $[a', b']$ is the i-th functional leftmost pseudo-zero.

See Figure 5.1 for a definition of ZERO1. Here $f([a, b])$ denotes the result of evaluating E with B_j substituted for variable x_j for all $j \in \{1, \ldots, n\} \setminus \{i\}$ and $[a, b]$ substituted for x_i.

5.2 Computing canonical leftmost pseudo-solutions

The function ZERO2 is defined recursively to compute the relational leftmost pseudo-zero; see Figure 5.2.

It assumes and maintains the invariant that when ZERO2 is called with a and b as arguments, the interval $[a, b]$ is an i-th leftmost pseudo-solution. The result is $\emptyset$ when it has been proved that no solution exists in $[a, b]$. If the result is not $\emptyset$, then it is $[a_0, b_0]$ such that $[a_0, b_0] \subseteq [a, b]$ and $a_0 \leq b_0$, and $[a_0, b_0]$ is the i-th canonical leftmost pseudo-solution.

Definition 18 (Functionally truncated probing) *Let $\mathcal{C}$ be the constraint system associated with an equation system.* Functionally truncated probing *of $\mathcal{C}$ with a constraint A refers to the result (failure or nonfailure) of an initial functional segment of a trace of the constraint system $\mathcal{C}'$ that results from adding A to $\mathcal{C}$. The result is nonfailure if* 0 *is contained in all intervals of variables occurring in the equality constraints in $\mathcal{C}'$; failure otherwise.*

If in ZERO2 one would replace probing with functionally truncated probing, then an algorithm would result that is for practical purposes equivalent to ZERO1. That is, functional box consistency can be computed with the same efficiency by constraint propagation on a system of atomic constraints, provided that propagation is not chaotic, but suitably structured and truncated.

6 Conclusions

Following [2, 7] we have treated systems of equations. Functional box consistency is easy to generalize to systems containing both equalities and inequalities. Relational box consistency seems more general because it applies to all interval constraint systems, not just to those that are derived from systems of equalities and inequalities.

In [2] `Newton` was compared with CLP(Intervals) on the Broyden Banded function. CLP(Intervals) was observed to required time exponential in the number of variables, whereas `Newton` required linear time. This is indeed to be expected: CLP(Intervals) used for pruning a single application of GCI. Because of the weakness of such pruning, the search tree reaches a significant depth. The size of such a tree is exponential in the number of dimensions.

In [2], the observation was made that `Newton` requires no branching on this example. Hence no exponential behaviour is to be expected.

It is now time to look beyond this particular example to those where even with pruning as powerful as in `Newton`, substantial branching is necessary. `Newton` showed that more effort spent in pruning is rewarded by a reduction in branching in such a way that the total computation time is much reduced. Thus there is a trade-off between time spent on pruning and time spent on branching: at some point, additional effort spent on pruning must stop being productive. This might suggest replacing functional by relational box consistency.

There is a better method. Note that to compute both functional and relational box consistency, one iterates all the way down to canonical intervals, the narrowest that the floating-point hardware allows. This is done to make the interval for one variable as narrow as possible. Yet to compute box consistency one has to do this for all variables repeatedly until no further narrowing is possible for any variable. In the beginning, most the intervals for most of the variables are still wide. While this is the case, it seems wasteful to iterate in functions ZERO1 or ZERO2 all the way down to canonical intervals: the convergence criterion should be adapted to the width of the other intervals.

In addition to this improvement, which applies both to ZERO1 and to ZERO2, there is an improvement that applies to the latter alone. As we described probing here, propagation is completed to convergence. As shown in Proposition 4, truncating the trace in propagation to the initial functional segment, causes ZERO2 to compute functional box consistency. By truncating the trace less drastically, say, till after the first cycle of a two-phase iteration, one obtains

better chance at getting failure in probing, yet avoids the negligible reductions associated with the last phases of propagation to convergence. It seems worth investigating how many phases are optimal in this respect.

7 Related work

The routine `absolve` in BNR Prolog uses probing to find narrower intervals than a single application of GCI can give. The mechanism was discovered independently by Chen and van Emden [5], who called it "hypernarrowing". It uses bisection to determine the greatest m such that probing with $x \leq m$ gives failure. In [5] "hypernarrowing" was used in optimization. A dramatic decrease in the number of function evaluations resulted. [5] missed the connection between "hypernarrowing" and box consistency.

Benhamou et al. [2] noted the ineffectiveness of the CLP(Intervals) `solve` routine for nonlinear equations. In response they introduced box consistency and used it to achieve dramatic improvements over the CLP(Intervals) `solve`. In the version they introduced (here called functional box consistency), they implicitly discard constraint propagation, and only use interval arithmetic.

The `HC4` algorithm of [3] is a propagation algorithm where instead of individual contractions one applies an algorithm called `HCRrevise`, which similar to a two-phase iteration truncated after the first cycle.

References

[1] K.R. Apt. The essence of constraint propagation. *Theoretical Computer Science*, 221(1-2):179–210, 1999.

[2] F. Benhamou, D. McAllester, and P. Van Hentenryck. CLP(Intervals) revisited. In *Logic Programming: Proc. 1994 International Symposium*, pages 124–138, 1994.

[3] Frédéric Benhamou, Frédéric Goualard, Laurent Granvilliers, and Jean-François Puget. Revising hull and box consistency. In *Proceedings of the 16th International Conference on Logic Programming*, pages 230–244. MIT Press, 1999.

[4] BNR. BNR Prolog user guide and reference manual. Version 3.1 for Macintosh, 1988.

[5] H.M. Chen and M.H. van Emden. Global optimization with Hypernarrowing. In *Proceedings 1997 SIAM Annual Meeting*, Stanford, California, 1997.

[6] M. Dincbas, P. Van Hentenryck, H. Simonis, A. Aggoun, T. Graf, and F. Berthier. The constraint programming language CHIP. In *Proc. Int. Conf. on Fifth Generation Computer Systems*, 1988.

[7] Pascal Van Hentenryck, Laurent Michel, and Yves Deville. *Numerica: A Modeling Language for Global Optimization*. MIT Press, 1997.

[8] Ramon E. Moore. *Interval Analysis*. Prentice-Hall, 1966.

[9] Arnold Neumaier. *Interval Methods for Systems of Equations*. Cambridge University Press, 1990.

[10] M.H. van Emden. Value constraints in the CLP Scheme. *Constraints*, 2:163–183, 1997.